



DTIC FILE COPY

INFORMATION SCIENCES INSTITUTE

UNIVERSITY OF SOUTHERN CALIFORNIA



4676 Admiralty Way Suite 1001  
Marina del Rey, California 90292-6695  
213/822-1511

AD-A231 025

# A RESEARCH PROGRAM IN COMPUTER TECHNOLOGY

## Final Technical Report

July 1981 - November 1989

Principal Investigator and Executive Director:  
Herbert Schorr

Prepared for the Defense Advanced Research Projects Agency  
Effective date of contract 1 July 1981  
Contract expiration date 30 November 1989  
Contract No. MDA903-81-C-0335  
ARPA Order 4242

DTIC  
ELECTE  
JAN 10 1991  
S E D

*This research is supported by the Defense Advanced Research Projects Agency under Contract No. MDA903-81-C-0335. Views and conclusions contained in this report are the authors' and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any person or agency connected with them.*

### DISTRIBUTION STATEMENT A

Approved for public release;  
Distribution Unlimited

## REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT This document is approved for public release; distribution is unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S)			5. MONITORING ORGANIZATION REPORT NUMBER(S) -----		
6a. NAME OF PERFORMING ORGANIZATION USC/Information Sciences Inst.		6b. OFFICE SYMBOL (if applicable)		7a. NAME OF MONITORING ORGANIZATION -----	
6c. ADDRESS (City, State, and ZIP Code) 4676 Admiralty Way, Suite 1001 Marina del Rey, CA. 90292			7b. ADDRESS (City, State, and ZIP Code) -----		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA/ISTO		8b. OFFICE SYMBOL (if applicable)		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA903-81-C-0335	
8c. ADDRESS (City, State, and ZIP Code) 1400 Wilson Boulevard Arlington, VA. 22209			1. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO. -----	PROJECT NO. -----	TASK NO. -----
			WORK UNIT ACCESSION NO. -----		
11. TITLE (Include Security Classification) Final Technical Report A Research Program in Computer Technology (Unclassified)					
12. PERSONAL AUTHOR(S) ISI Research Staff					
13a. TYPE OF REPORT Final Technical		13b. TIME COVERED FROM 7/81 TO 11/89		14. DATE OF REPORT (Year, Month, Day) 90 12 31	
15. PAGE COUNT					
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
09	02				
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
<p>This report summarizes the research performed by USC/Information Sciences Institute for the Defense Advance Research Projects Agency. The research is focused on the development of computer science and technology, which is expected to have a high DOD/military impact.</p>					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified		
22a. NAME OF RESPONSIBLE INDIVIDUAL Dan Pederson			22b. TELEPHONE (Include Area Code) 213/822-1511		22c. OFFICE SYMBOL

# Table of Contents

## 1. COMPUTER RESEARCH SUPPORT

- 1.1 BACKGROUND
- 1.2 PROBLEMS BEING SOLVED
- 1.3 GOALS AND APPROACHES
- 1.4 PROGRESS
- 1.5 MILITARY IMPACT
- 1.6 FUTURE WORK

0  
0  
1  
1  
3  
4  
4

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



# 1. COMPUTER RESEARCH SUPPORT

**Director:** Dan Pederson

## **Technical Staff:**

### **Software:**

Ray Bates  
Jim Koda  
William Moore  
Koji Okazaki  
Tom Wisniewski  
Craig Ward  
Dwight Fromm

### **Hardware:**

Ray Mason  
Glen Gauthier  
John Patrick  
John Scofield  
Ramon Gonzalez  
Fred Grolle  
Val Fucich

### **Operations and Network Services:**

Vicki Gordon  
James Hurd  
Roger Lewis  
Wayne Tanner  
Mike Zonfrillo  
Joe Kemp  
Chris Refuerzo  
Sean Schur

## **Support Staff:**

Manon Levenberg  
Pat Thompson

## 1.1 BACKGROUND

ISI provides cost-effective and key computer support to researchers at ISI and DARPA, DARPA contractors and affiliates, and the military services. In addition to providing raw computing service on TOPS-20 and Unix to members of the DARPA research community and DARPA itself, ISI also originally developed and provided and/or now helps to support many additional mature software packages for the entire Internet community, including text editors (XED, Gnu Emacs), mail handlers (Hermes, MSG, SNDMSG, MM), compilers (PASCAL, C, FORTRAN, COBOL, MACRO, BLISS), language environments (Interlisp, Mainsail, Ada), and network access tools (Telnet, FTP, SMTP).

ISI has built a reputation for excellence and efficiency in both experimental computer services and production services, and has repeatedly demonstrated its continuing high standards to users while maintaining a relatively small expert staff. ISI provides guidance to DARPA on what constitutes a sensible load for a TOPS-20 host, and also provides management control to ensure an adequate level of support for the DARPA user community.

ISI has also shown itself to be extremely proficient in the realms of network access and host security, successfully walking the fine line between protecting the interests of the community and having that security be an encumbrance to the user community. Security involves a number of access controls and, perhaps more important, it includes protection schemas, password encryption checkers, and monitoring demons that can be activated when a problem is suspected.

ISI and its current research projects have benefited substantially from in-house competence in computer

services, particularly as the original support of the TOPS-20/DEC-KL environment has widened to include an eclectic and complex collection of hardware and software architectures. The staff of the Computer Center has grown in expertise and maturity to meet the demands of this wider collection of equipment and software, while the overall size of the group has been reduced.

With the continuing evolution of computer workstation technology, ISI expects to see a decline in the demand for network-based, interactive, timesharing cycles as users move into workstation local area network environments.

## 1.2 PROBLEMS BEING SOLVED

The Computer Research Support project is responsible for providing reliable computing facilities on a 24-hour, 7-day schedule to the Internet research and development community. At the same time, the project makes available to Internet users the most current releases of hardware and software on the supported machines. The project provides continuous computer center supervision and operation, and a full-time customer-service staff that responds to user inquiries. This project supports a major computer installation at ISI's main facility in Marina del Rey, California. Shared staff within the facility lends infrastructure support to ISI's internal research efforts

## 1.3 GOALS AND APPROACHES

The ISI Information Processing Center provides support in four distinct, though tightly interrelated, areas: Hardware, Systems Software, Operations, and Network Services. The overall charter of the organization is to assure that the needs of the user community are addressed and protected as efficaciously as possible. To achieve this end, each group is concerned about the effective use of the machines and software tools, and about the security of the physical plant, the system files, and other online information. The more specific goals and approaches of each group are summarized below.

### Hardware

To achieve a reliability goal of 98.7 percent scheduled uptime, preventive and remedial maintenance responsibilities have been assigned to an in-house computer maintenance group. This group provides cost-effective 24-hour, 7-day coverage. To maintain the reliability goals, preventive maintenance is very closely controlled, and online diagnostics and analysis are emphasized. A primary component in the reliability and availability of the hardware is the physical environment in the computer facility itself. Accordingly, significant time and resources are expended in ensuring that the best, most cost-effective environmental controls are at the facility.

## **System Software**

The software group's overall goal is to install and maintain, at maximum reliability, ISI's VMS, UNIX, and TOPS-20 operating systems and applications software. In order to accomplish this goal, the group provides 24-hour, 7-day coverage to analyze system crashes and to provide appropriate fixes. In addition, it is the group's responsibility to install, debug, and modify the latest monitor and kernel versions, and the associated subsystems, available from the vendors.

## **Operations**

The Operations staff is responsible for monitoring the systems and the physical environment of the computer room itself. The facility has at least one operator on site to monitor the facility which runs 24-hour, 7-days a week. This goal is achieved through a variety of means, including regularly scheduled full and incremental backups of all systems; permanent archivals of requested or infrequently accessed system and user files; magnetic tape storage and pointers to all information extant at the time of removal of directories from the various systems; and, perhaps most important, redundant off-site storage of all significant information active on the disk structures or existing on tape within the facility.

When a problem occurs, the on-duty staff isolates it and takes appropriate action. They work closely with the on-call hardware and software support staff in resolving system problems so that maximum uptime can be achieved. On the night and weekend shifts, the Operations staff responds directly to user inquiries. Proper training, experience, continuity, and familiarity with the environment are especially stressed.

## **Network Services**

Network Services, the ISI customer-service group, provides a two-way communication link between the users and the rest of the support staff. The group handles all directory management issues, monitors available disk space, answers questions about the primary software packages, solves a number of lower level technical questions, and will also cover the Operations duties when required. This support is accomplished by maintaining a coverage from 0600pst to 1800pst Monday thru Friday, for prompt problem resolution and rapid information exchange, both on-line and by telephone. The group also offers introductory training in the use of hardware and software tools available on the ISI systems, as well as providing documentation for new users of the Internet. Network Services also assists in the formulation of user training programs for large, Internet-based military experiments at, for example, the Strategic Air Command, Offutt Air Force Base, Nebraska; the Naval Postgraduate School, Monterey, California; and the Systems Design Center at Gunter Air Force Base, Alabama.

Appropriate documentation is constantly being generated and distributed to individual users, as well as to remote user-group liaison personnel; this documentation ranges from simple, beginner-level explanations to more technical information suitable for experienced users. In accordance with ISTO guidelines, the customer-service group provides regular system utilization accounting data to DARPA.

## 1.4 PROGRESS

The ISI Computer Center Facility operated without major hardware or software difficulties during the reporting period. Scheduled uptimes averaged well above 99 percent. Three of the four DEC 2060 computers running the TOPS-20 operating system was shut down in 1988. The development of plans to move users and then shut down of the last DEC 2060 hosts also began during this period. The users of these system will be transferred to a SunServer running BSD 4.3 UNIX.

### Hardware Additions

The ISI Computer Center continued its move away from DEC 2060 computers toward personal workstations and file servers. The facility acquired a large number of workstations and servers.

The hardware staff performed upgrades and additions to the hardware configurations on the DEC 2060 (including security improvements). An additional change included more memory on ISLA. Most of the old RP06 disk drives were phased out in favor of larger, more modern and cost-effective Winchester technology RP07 drives. The RP06 drives that remain are now mostly used for mountable structures. All installations and changes were carried out during scheduled evening maintenance hours without interfering with normal operations.

The Computer Center houses and or lends substantial support to the array of personal computers (PC's), workstations, and symbolic processing engines acquired over the last several years by a variety of projects. The Center assures that the filesystems of these various machines are backed up and that network connectivity (where required) is robust.

The current list of major multiuser processors, network servers, and supported individual machines (PC's, workstations, and symbolic processing computers) follows:

10 DEC VAX 11/750 computers	1 Micom data switch (2 bays)
3 DEC VAX 11/780 computers	1 LeeMah data security switch
2 DEC VAX 8650 computers	1 Xerox 5400L Penguin laser printer
16 DEC MicroVax computers	1 Dataproducts high-speed printer
4 Imagen laser printers	2 Xerox 2700 laser printers
18 Symbolics 3600/3645/3675 LISP machines	1 Xerox 8045 print server
65 IBM-PCs (various models)	2 Xerox 8031 file servers
58 SUN Microsystems workstations	40 desktop printers
20 Hewlett-Packard Bobcat workstations	3 BBN Butterfly gateways
17 Texas Instruments Explorer workstations	2 Iris workstations
10 Apple IIE/Lisa/Macintosh computers	120 modems (300, 1200, 2400 Baud)
20 computer terminals (various models)	1 Connection Machine
1 Symult Parallel Processor	

## **System Software Enhancements**

Updating of the mail handlers and the Domain name space were the major software enhancements. Other software that was upgraded includes: Scribe, and Gnu Emacs.

## **1.5 MILITARY IMPACT**

ISI is perhaps the finest university-based research center promoting the sharing of software resources within the DARPA community; it assumes responsibility for providing support to key DoD community personnel so as to demonstrate the great utility of the ARPANET and MILNET resources. The effective and rapid transfer of information, electronic mail and other data, computer programs and tools illustrates, on a working daily basis, ways to enhance military efficiency. Specific technology transfer of relevant research to the military is heavily dependent on the facility.

ISI's Computer Center provides ARPANET, MILNET cycles and support 24 hours a day, 7 days a week to the Strategic Air Command, Naval Postgraduate School, Gunter Air Force Base, the Office of Naval Research, and the ADDS Experimental Test Division at Fort Bragg, North Carolina, as well as to the contractors, researchers, and administrators of militarily significant research coordinated out of the DARPA office in Washington, D.C. In addition to supplying machine time and directed Network Services support, this project continues to provide substantial additional support in the following areas:

- General accounting information as needed for workload analysis on the various machines.
- Rapid response to user requests and problem reports.
- Tailored security awareness, and manual and automated tracking.
- Maintenance and upgrading of electronic mail system software, shared online bulletin boards, and other specialized communications and file-transfer programs.
- Maintenance of approximately 1500 user directories (as well as hundreds of support and overhead directories) on the TOPS-20 machines used by the DoD and affiliates.

## **1.6 FUTURE WORK**

The Computer Research Support project will continue to provide computing service to the DARPA research community, provide and support software packages for the Internet community, and offer a program of technology transfer and user education through the Network Services group.

Some specific planned activities for the next year include the following:

- Complete the subnetting at ISI thus improving the computer response time.
- Installation of new procedures and support for acquisitions of equipment at ISI. This will include potential new products (either purchased at DARPA's direction or via grants) from the following vendors: Hewlett-Packard, Texas Instruments, SUN Microsystems, Symbolics, Digital Equipment Corporation, and others



- Installation of new releases of operating systems on all supported computers.
- Upgrading disk storage to a newer more cost effective technology.
- Upgrading tape storage to a newer more cost effective technology.

# Table of Contents

<b>1. DARPA HEADQUARTERS SUPPORT CENTER</b>	<b>0</b>
1.1 PROBLEM BEING SOLVED	0
1.2 GOALS AND APPROACH	0
1.3 SCIENTIFIC PROGRESS	0
1.4 MILITARY IMPACT	1
1.5 FUTURE WORK	1

# **1. DARPA HEADQUARTERS SUPPORT CENTER**

## ***Research Staff:***

Dan Pederson

Ray Mason

John Reed

## **1.1 PROBLEM BEING SOLVED**

The ISTO computer center environment requires sophisticated systems capable of providing services for program managers and analysts commensurate with the complexity of their tasks. The systems must be supported in a cost-effective manner, while maintaining a less than two percent downtime.

## **1.2 GOALS AND APPROACH**

To provide cost-effective service, ISI devised a set of remote environment surveillance and system diagnosis tools. These include remote monitoring (from ISI) of the ISTO computer room's temperature, humidity, power, fire-suppression system, VAX system crashes, and computer room entry via an automatic reporting system that uses standard telephone lines. ISI personnel are able to observe the status of the ISTO computer center through a series of video cameras whose images are transmitted to ISI via the ARPANET. Access to the Sun server systems for system software maintenance and enhancement is provided via remote dial-in. Hardware problem diagnosis is also achieved via remote dial-in, where diagnostics can be run from ISI. ISI hardware technical staff can, upon analysis of the diagnostics, determine the failing module and with the help of selected ISTO personnel, perform module swaps.

## **1.3 SCIENTIFIC PROGRESS**

During this period two Sun 3/180 systems were added to the ISTO computer room; one for the Oracle database and one to handle 19.2KBPS communications between the home Suns and the ISTO servers.

To allow ISI personnel remote access to the VAXes and Sun servers, ISI installed additional phone lines, an 800 service, and additional modems.

Dual porting on the Sun disk drives to enable another server access if the primary CPU is down for an extended period was implemented. The switch is transparent to the user.

The revamping of the ISTO Ethernet to a "thinnet" schema to provide ISTO with a more maintainable and expandable network has been completed. From the time of its completion ISTO has experienced no ethernet outages.

Scheduled maintenance of all equipment was accomplished.

## **1.4 MILITARY IMPACT**

ISI has proven that remote maintenance for small computer facilities is feasible and cost effective. ISI has been able to maintain a high mean time between failures for the existing systems through effective electronic and video surveillance of the computer room environment and remote access to the VAX and Sun systems in the computer room for maintenance and diagnostic purposes.

The "remote" approach to computer management of small unattended computer facilities could be a cost-effective solution for military applications.

## **1.5 FUTURE WORK**

Remote surveillance of the computer facility for maintenance and security purposes will continue. ISI is investigating the possibility of upgrading the remote video capability to allow ISI to receive video images closer to real-time. The scan compress transmit time causes a delay, which could be partially alleviated by a new design of the compression algorithm, although much of the delay is caused by network load.

The full implementation of the LeeMah dial-back security system is ready. The final phase of this implementation has been put on hold pending a detailed system security review.

As wider and more complex usage of the Sun workstations takes place, additional disk storage will be required. This additional storage will also entail optical storage system(s). Also, existing workstations will be upgraded to provide the latest computing environment to program managers.

One major task will be planning the computer facility move to its new location, which is scheduled for late 1990 or early 1991.

# Table of Contents

<b>1. STRATEGIC COMPUTING DEVELOPMENT SYSTEMS</b>	<b>0</b>
1.1 PROBLEM BEING SOLVED	0
1.2 GOALS AND APPROACH	0
1.3 SCIENTIFIC PROGRESS	2
1.4 IMPACT	2
1.5 FUTURE WORK	2

# 1. STRATEGIC COMPUTING DEVELOPMENT SYSTEMS

## *Research Staff:*

Dan Pederson

Ray Bates

Ray Mason

## 1.1 PROBLEM BEING SOLVED

Using recent advances in workstation/server architecture, artificial intelligence, and computer science, DARPA plans to create a new generation of "machine intelligence technology." The DARPA Strategic Computing program will be targeting key areas where advances can be leveraged towards broad advances in machine intelligence technology and the demonstration of applications of the technology to critical problems in defense.

The Strategic Computing program will be supported by a technology infrastructure. This infrastructure will serve to bootstrap the program by providing current state-of-the-art computing technology, network communications, and shared resources to the program participants.

The aim of this project is to provide development computers for the Strategic Computing program. System integration and the distribution of the systems to the program participants, as well as a defined architecture for system communications and resource sharing among the computers acquired, are requisite for this project.

One of the project's aims is to provide cost-effective and state-of-the-art engineering workstations to the Strategic Computing community. As the research in this program has progressed, and as the sophistication and complexity of available hardware and software in the vendor community have evolved, the more generalized workstations have proven to be a valuable adjunct to the already established base of specialized symbolic processors.

## 1.2 GOALS AND APPROACH

A number of machines have been developed to support high-speed processing of large symbolic programs. Each of these machines provides an extensive interactive programming environment, a sophisticated display-manager "window system," a real-time, window-oriented editor, incremental compilers, and dynamic linking.

These systems are the state of the art in program development environments. They are widely used in the research community for systems development, expert systems, natural language systems, mapping applications, and support of CAD/CAM environments.

Since many of these are single-user systems, they cannot be timeshared in the traditional sense and thus the cost per researcher is high. To bring these costs down, workers are currently placed in the awkward situation of having to schedule their computing needs. The resultant scheduling conflicts naturally lead to low researcher productivity.

An examination of dynamic machine use shows that many activities do not require the high-speed processing capabilities afforded by these machines. Less expensive, general-purpose machines supporting the same language base have become available only over the last few years. While formerly adequate only for less-intensive research activities such as editing, text preparation, and file scanning, these general-purpose machines are now sufficient for more intensive program development and execution activities.

The duality of resources required for different types of activities and the availability of machines of appropriate cost/performance for carrying out those activities suggests a workstation server architecture based on these two classes of machines and an interconnecting network. In order to support dynamic source code exchange between server and workstation, it is necessary for each to support the same AI language system. Common LISP is the natural choice in such an architecture, as it was designed with the concepts of commonality and portability as primary goals.

A significant fallout of such an architecture, when viewed as workstations and servers on TCP/IP-based networks, is the ability of researchers to communicate over the Internet in order to use high-powered resources available on prototype and low-production machines. This has been particularly useful on prototype machines being developed under the machine architecture phase of the program.

ISI has acquired a mix of these machines to support the requirements of the Strategic Computing program, and has helped to integrate and test individual systems as required for particular Strategic Computing program participants.

The integration tasks have avoided duplication of effort among Strategic Computing participants. ISI has collected software to support this architecture as it has been developed by vendors and the research community. ISI does a modest amount of testing, modifying, and augmentation of software; provides support for software exchange and distribution among DARPA development sites; helps assure working compatibility of the Common LISP systems; works with commercial vendors to resolve vendor-specific problems; and works to allow network compatibility of the systems with DoD network protocols (TCP/IP). The software and system integration efforts are carried out at the ISI Marina del Rey facility.

### 1.3 SCIENTIFIC PROGRESS

During the reporting period, additional general-purpose configurations were acquired. Some were installed and integrated at ISI, while most were shipped out to research facilities, educational sites, and government agencies.

In the past year, 85 machines were acquired as part of this Strategic Computing project, as well as a large variety of component peripherals, software packages (from Interleaf, CCA, Unipress, etc.), and software maintenance agreements. The major machines acquired were:

- 20 - Apple Computers
- 10 - NeXT Computers
- 40 - Sun Computers
- 15 - PCs, various models
- Various hardware peripherals packages

### 1.4 IMPACT

The DARPA Strategic Computing program will continue to develop and integrate advanced computer technology into military applications. Technological advancement will be sought in high-speed symbolic machines, as well as application of this technology to the military. Potential military application of this technology includes autonomous systems (land, air, and sea), battlefield management and assessment, planning and advising, and simulation systems.

The initial program applications included an autonomous land vehicle, a pilot's associate, and a carrier battle group battle management system. These applications test the evolving technology in the areas of vision, expert systems, speech recognition, and high-performance knowledge-processing systems. Each application seeks to demonstrate the new technologies potential for providing a major increase in defense capability and to reflect a broad demonstration base that is relevant to each of the three services.

The development systems acquired through this project will support the technology base and the targeted applications under the Strategic Computing program.

### 1.5 FUTURE WORK

Several manufacturers have perceived that there will be a substantial marketplace for a more general-purpose symbolic processing engine in the future. We expect that the price performance ratio of these new machines will continue to improve at a rapid pace, resulting in more efficient use of researcher time as more state-of-the-art machines can be acquired at the most reasonable price.



During the next year, ISI will continue to negotiate with vendors and acquire a mix of high-end LISP machines and other engineering workstations capable of handling symbolic processing in addition to other applications. As the workstation/server architecture is defined by DARPA, ISI will configure systems, test vendor software, and distribute the systems to the Strategic Computing program participants.

# Table of Contents

<b>1. NEW COMPUTING ENVIRONMENT</b>	<b>0</b>
1.1 PROBLEM BEING SOLVED	0
1.2 GOALS AND APPROACH	0
1.3 SCIENTIFIC PROGRESS	1
1.4 IMPACT	2
1.5 FUTURE WORK	3

# 1. NEW COMPUTING ENVIRONMENT

## *Research Staff:*

Dale Chase  
Jim Koda

## *Support Staff:*

Glen Gauthier  
Manon Levenberg  
Ray Mason

## 1.1 PROBLEM BEING SOLVED

For the past decade, the computing needs of computer science researchers have been provided on large-scale, timesharing machines. These researchers no longer find these timesharing machines a viable approach for their research. The high-speed and dedicated workstation has become the machine of choice. ISI's DEC mainframe running TOPS-20 is at full capacity and become very costly to operate. Additionally, it is not capable of supporting many of the current and proposed research efforts at ISI.

As techniques and expectations (especially in the area of Artificial Intelligence) have advanced, general-purpose machines are no longer able to provide the style of interaction and capabilities, in terms of throughput and address space, that are required to support current research efforts. The NCE project continues to:

- provide a very significant improvement in both the amount and the quality of available computing resources to the ongoing ISI research efforts through the use of dedicated personal workstations and higher capacity centralized processors and servers.
- provide for diverse access methods to the common set of services.
- fully support the use of special-purpose workstations and processors as required by individual research projects.
- free up capacity on the existing mainframes by offloading some common functions to central servers.

## 1.2 GOALS AND APPROACH

The ISI research community is made up of a diverse set of projects and cultures. In order to best support a community with such varied needs, we decided to implement a heterogeneous environment consisting of a variety of personal workstations, dedicated server, and continued use of our existing timesharing systems. The New Computing Environment was designed to consist of local processing nodes (workstations and multiuser mainframes) connected via a local network to a set of central servers. A certain minimum set of functions must be available to all nodes in this environment: mail, file manipulation (access and transfer), and Internet connectivity (Telnet). The environment provided by these nodes and servers conceals the fine-grain detail from outside users; users external to the environment need not know what workstation a particular user is on, unless they want to. At minimum, these services

should be available via any of several communication media: local area network, Internet, and dialup access. The NCE provides additional support as needed to allow low-end workstations on our mainframes to use these services as easily as the fully capable workstations in the Sun, Silicon Graphics, DEC VAXstations and Apollo class. Primary access to these services is via the DoD TCP/IP protocol suite, with some extensions and enhancements to support particular modes of interaction that are not yet supported (e.g., random file access and remote procedure call).

The use of these services by our mainframes has resulted in immediate improvements for all users, as many CPU-intensive activities have been offloaded to servers. Making these services available from our mainframes also provides for their use from home systems. Access for these devices is via dialup lines to our existing mainframes, or to the servers themselves.

The following services are provided to all nodes in the environment:

- centralized mail service
- centralized file service
- document formatting
- high-quality printing
- communications
- specialized processing

While the principal incentive for the New Computing Environment was the support and integration of a diverse set of dedicated personal workstations, not everyone at ISI requires this kind of hardware and capability. Until there is an actual need to migrate research efforts to workstations, there will be a significant population content with our existing VAX mainframes. However, the enhanced facilities offered by the provision of centralized servers also benefits this population. These enhancements take the form of increased accessibility of files and mail, increased reliability in terms of both data integrity and security, and improved performance of the existing mainframes due to the offloading of CPU-intensive activities to the dedicated servers.

To reduce the operating and maintenance costs of in-house computers, we have moved groups of users from TOPS-20 to VAXes running both UNIX and VMS.

### **1.3 SCIENTIFIC PROGRESS**

The main computing resource at ISI, a DEC VAX 8850, was updated to run Berkeley UNIX, version 4.3. We installed a number of Sun file servers on multiple subnets to improve system performance. This same architecture and hardware was transferred to the DARPA ISTO environment and eliminated daily outages.

The initial design and development of a system that makes existing archived files available over the Ethernet to any other system was completed and successfully installed. The Massbus Ethernet Interface System (MEIS) on A.ISI.EDU was enhanced to provide better performance to our remote users since the ARPANet was being shut-down.

Enhancements to the root domain server for the Internet were completed. Work was done to integrate the domain server code into the 6.1 environment. A side benefit is that the A.ISI.EDU user community has been ushered into the domain name world, instead of being restricted to communicating with hosts in a static host table.

In other areas, work has been ongoing to improve the connectivity between ARPANET, MILNET, and ISINET systems. Different systems on the ISINET have been configured and tested to act as private gateways between the ISINET and the MILNET, and dramatic improvements have been observed. Routing strategies designed to make the most of such gateways were tested and installed.

Several improvements were made to the PC integration into the ISINET. New versions of Sun PC NFS were installed on local IBM PCs. A new release of the CMU PC IP was obtained and installed on local PCs, employing a network-based installation scheme called Bootp to allow previously networked PCs to be updated without the need to shuffle floppies. Local improvements to the PC IP code were folded in, including an enhanced lpr client that allows files to be spooled from PCs to a UNIX server for printing to any of the printers in the ISI environment. Various ethernet interfaces were tested on portable PCs, software tested, and delivered to DARPA ISTO.

## **1.4 IMPACT**

Our experience with personal computers and workstations has made it clear that they represent a useful alternative to large, timeshared mainframes for certain research activities. Projects within the institute have moved their work to workstations, and have therefore been able to continue research that exceeded the capacity of our mainframes.

As the reliability of added facilities and of the entire environment is proven at ISI, we offer the same type of environment to outside contractors via the Exportable Workstation Systems project. These two projects combine to provide an ideal environment for developing and refining the facilities and capabilities that are becoming increasingly important in the command and control context. These capabilities include the handling of redundant databases and the support of a heterogeneous collection of hardware.

## 1.5 FUTURE WORK

We will study subnet configurations of the local area network to optimize network traffic. To aid in this study, we plan to implement an extended network traffic analyzer that will provide detailed short-term logs and long-term traffic histograms.

The Laser Disk Archive project will be completed and installed at ISTO. We plan to study approaches that will allow remote workstations to gain access to network-based resources via dial-up lines and high-speed dedicated links.

We will also specify and implement a full-function mail service for IBM-PC's. The service will be integrated into the local area network and will interface to Internet mail services.

# Table of Contents

<b>1. COMMERCIAL MAIL</b>	<b>0</b>
1.1 PROBLEM BEING SOLVED	0
1.2 GOALS AND APPROACH	0
1.3 SCIENTIFIC PROGRESS	1
1.4 IMPACT	2
1.5 FUTURE WORK	2

# 1. COMMERCIAL MAIL

## *Research Staff:*

Dale Chase

Craig Ward

## *Support Staff:*

Manon Levenberg

Glen Gauthier

## 1.1 PROBLEM BEING SOLVED

The evolution of large electronic mail systems testifies to the increasing importance of electronic mail as a means of communication and coordination throughout the scientific research community. These systems include the DARPA Internet mail system, the GTE Telemail system, the MCI-Mail system, and the IEEF Compmail system on ITT Dialcom. Until now these systems have operated autonomously, and no convenient mechanism has existed to allow users of one system to send electronic mail to users on another system. The Intermail system, developed by the Internet Concepts Research project at ISI, demonstrated a mail-forwarding system that allows users to send electronic mail across such mail system boundaries. The Commercial Mail project will convert this system from a research project into a commercial product.

## 1.2 GOALS AND APPROACH

The most significant limitation of the Intermail system is its inability to handle forwarding to more than one "other" mail system in a single interaction--that is, each message can be delivered to only one other mail system. We will eliminate this limitation by introducing a new syntax for addresses. Any number of addressees may be specified on any number of "other" mail systems. This will allow users on "other" mail systems to be included in mailing lists along with DARPA Internet recipients, UUCP recipients, CSNET recipients, etc.

It is presently possible for DARPA Internet users to communicate easily with users of both UUCP and CSNET via cooperating hosts that maintain connections with these communities. Addresses are specified in the local syntax of each of the three systems, and conversions are handled by the forwarding hosts. This is the way we intend to have the commercial mail forwarding facility function. This is a relatively easy facility to provide on the DARPA side, where the address syntax is well understood and flexible, and the entire process will be under our control. On the commercial side, we will continue to use the techniques developed by the Internet Concepts project, which use forwarding information embedded in the text portion of the message in a Simple Forwarding Header (SFH). Our conversion system will, however, make it possible for DARPA Mail recipients to reply directly to messages from commercial users using existing DARPA Mail composition programs.



### 1.3 SCIENTIFIC PROGRESS

Upon the successful development of the Intermail system, the ISI computer center software group was contracted to build a robust, production-quality, operating Commercial Mail Relay (CMR) system. In its first year the CMR project spent a great amount of time evaluating the Intermail research project software, various operating system platforms, and network mail system software suitable for building the foundation of a commercial product. With the selection of 4.3 bsd UNIX for the operating system and the Multi-channel Memo Distribution Facility (MMDF) as the system mailer, the first specifications for the CMR were written.

UNIX and MMDF were selected for system software so that the CMR could be developed in a highly portable environment. Both UNIX and MMDF run on a variety of platforms. MMDF, with its notion of a dedicated channel for messages going to a particular destination, allows great flexibility concerning the addition or removal of commercial networks.

Development began on a VAX/11-750 with the target production machine being a Microvax VAXstation II. During the development period, numerous demonstrations of the system were given, enabling researchers at ISI to give the CMR development team positive feedback. The first CMR MMDF channel and a pre-processor for SFHs were put into service in October 1987.

This channel bridges the GTE Telemail system to the Internet. It interfaces to the Telemail system by simulating a human using a terminal. Negotiations are under way with GTE to update their software to allow direct computer access. Because the GTE system assumes that a human is entering the system, mail errors (such as misaddressed mail) are handled in a verbose format. This makes error handling in CMR very challenging. The Telemail channel and the SFH pre-processor currently handle some failed mail situations, such as incorrect commercial mail system, incorrect commercial mail user name, and incorrect commercial mail host. We plan to include others and enhance the understanding of CMR-generated error messages as our experience with the CMR user community increases.

After the Telemail channel and SFH pre-processor stabilized, the CMR development team began designing a channel for ITT Dialcom. Several organizations, government and private, operate on Dialcom systems. The CMR Dialcom channel will be used by IEEE, ONR, and NSF.

While the design and coding of the Dialcom channel proceeded, bugs and other required changes were made to the Telemail channel and the SFH pre-processor, along with bug-fixes and upgrades for the system software. User documentation was updated and the program logic manual was edited to reflect changes and in response to user needs. Simple accounting programs were developed to track the usage.

Work continues on refining and expanding the Commercial Mail Relay.

## 1.4 IMPACT

The availability of high-performance personal workstations and dedicated special-purpose processors, and the preference for these machines by researchers, have made the acquisition and operation of large mainframes capable of providing reliable mail service unacceptably high. Since most of these research machines cannot provide mail service, the Commercial Mail implementation will allow the users of these machines to remain accessible to the rest of the research community without maintaining two distinct computing facilities.

The availability of a reliable high-capacity bridge between commercial mail systems and the DARPA Internet will provide better communication among contractors, especially those involved in the Strategic Computing Program, who might otherwise be denied convenient access to colleagues.

## 1.5 FUTURE WORK

ISI will continue to update and support the CMR system. We plan to add software to the system that will allow mail to be relayed from the Internet to Dialcom (for ONRMail, IEEE Compmail, and NSFMail) and MCI Mail. The CMR is flexible enough that the addition of other commercial mail networks such as Compuserve, GENIE, or AT&T Mail will be possible depending on the interest expressed by the Internet community in any such networks. ISI will also provide the following enhanced services:

- negotiate more cost-effective billing procedures from commercial networks
- negotiate the implementation of new machine interfaces with commercial mail networks
- improve systems reliability
- implement software that will allow faster file transfer
- improve systems accounting

The improved systems accounting software will include the ability to recharge costs to remote sites that use the CMR system.

Once this software is developed, ISI will serve as a distribution center for outside agencies that want to use this type of mail relay software. This software will be fully supported by a set of user documentation. ISI will maintain an electronic mailbox that can be used by remote users to report bugs, problems, etc. ISI will also distribute the CMR software to other sites that run UNIX.